

# Seven of Eight AI Deployment Patterns Are Structurally Suppressed

The 2026 landscape propagates on capital, distribution, and lock-in — not on structural merit.

THE GROVE FOUNDATION · AI PATTERN BENCHMARK · MARCH 2026 · CC BY 4.0

Date of Analysis: March 8, 2026

---

## What this benchmark measures

Gartner measures market momentum. Forrester measures vendor capability. IDC measures competitive positioning. None of them asks the question that determines whether a technology wins the 36-month race: does the pattern propagate on its own structural merits, or does it require continuous external subsidy to sustain adoption?

The  $\Lambda$  framework measures that question directly. It treats adoption as a structural mechanics problem. Capital can accelerate inferior architectures for three to five years. Distribution power can sustain them inside an installed base. Neither subsidy survives the displacement pressure that accumulates once cognitive friction drops and an alternative becomes structurally accessible.

Eight AI deployment patterns. One equation. Observable data only. Ninety-six cited sources. No vendor narratives.

## The finding

One pattern is Approaching Critical: Mistral/DeepSeek at  $\Lambda = 0.0314$ . Two patterns sit Sub-Critical. Five patterns — including all three platform-bundle and centralized-API incumbents — are Structurally Inert.

The \$650 billion AI infrastructure buildout is structurally suppressed. It propagates on venture capital, enterprise distribution power, and accumulated contractual lock-in. Remove the subsidy and the propagation stalls or reverses.

## The equation

$$\Lambda = (S \times R \times V) \cdot [ 1 / (1 + (\beta \cdot Fc)^2) ]$$

Spreadability (S), Standardized Rails (R), Validation Multiplier (V), Cognitive Friction (Fc), Exogenous Incentive ( $\beta$ ) — each a structural input grounded in observable data.

## Historical calibration

Case	$\Lambda$	Tier
TCP/IP (vs. OSI)	<b>0.452</b>	Critical Mass
ISO Shipping Container	<b>0.380</b>	Critical Mass
Bitcoin	<b>0.122</b>	Critical Mass
US Metric System	<b>0.003</b>	Structurally Inert

## APRIL STRUCTURAL SHIFT — Note on Gemma 4

The data and  $\Lambda$  scoring in this March 2026 benchmark were finalized prior to the early April release of Gemma 4 under an Apache license. Full mathematical impact will be reflected in the Q2 landscape update, scheduled for late June.

### The 2026 landscape, scored

Pattern	Category	S	R	V	Fc	$\beta$	$\Lambda$	Tier
<b>Mistral / DeepSeek</b>	Open-Weight	0.80	0.60	1.0	6.0	0.630	<b>0.0314</b>	Approaching Critical
<b>Meta Llama</b>	Open-Weight	0.70	0.70	1.0	5.0	1.357	<b>0.0104</b>	Sub-Critical
<b>Apple Intelligence</b>	On-Device	0.20	0.57	1.0	2.0	1.710	<b>0.0090</b>	Sub-Critical
<b>Anthropic Claude</b>	Centralized API	0.47	0.80	1.0	5.0	1.587	<b>0.0059</b>	Sub-Critical
<b>OpenAI GPT</b>	Centralized API	0.47	0.93	1.0	6.0	2.924	<b>0.0014</b>	Structurally Inert
<b>Google Gemini</b>	Platform Bundle	0.30	0.80	1.0	5.0	2.924	<b>0.0011</b>	Structurally Inert
<b>Microsoft Copilot</b>	Platform Bundle	0.30	0.80	1.0	8.0	5.000	<b>0.0002</b>	Structurally Inert
<b>Autonomaton</b>	Sovereign Open	0.93	0.83	0.2	6.0	6.300	<b>0.0001</b>	Structurally Inert

### Two structural clusters

**The Dependency Profile** — high rails, low spreadability. OpenAI, Google, Microsoft, Anthropic. Easy to adopt. Hard to leave.

**The Sovereignty Profile** — high spreadability, variable friction. Mistral/DeepSeek, Meta Llama, Apple Intelligence, Autonomaton. Harder to adopt. Easier to leave.

## Eight AI Adoption Patterns Scored

**OpenAI GPT** —  $\Lambda = 0.0014$ , **Structurally Inert**. Highest Standardized Rails in the landscape ( $R = 0.93$ ). Structurally Inert anyway. The GPT-4o February 2026 retirement broke production workflows at every enterprise tuned to that specific model. *Rails without sovereignty is a trap, not an advantage.*

**Anthropic Claude** —  $\Lambda = 0.0059$ , **Sub-Critical**. Lower cognitive friction via ISO 42001, structured safety, cleaner enterprise onboarding. The structural pattern is identical: rented, centralized dependency.

**Microsoft Copilot** —  $\Lambda = 0.0002$ , **Structurally Inert**. Lowest-scoring pattern. 3.3% paid conversion across 450M commercial M365 seats. A sales motion sustained by distribution power, not a structural pattern.

**Google Gemini** —  $\Lambda = 0.0011$ , **Structurally Inert**. Fragmentation (Gemini / Vertex AI / AI Studio / Agent Builder) compresses rails below what infrastructural strength would otherwise suggest.

**Apple Intelligence** —  $\Lambda = 0.0090$ , **Sub-Critical**. Lowest  $F_c$  in the landscape (2.0). Closed hardware crushes spreadability. The organization does not build transferable capability. It builds Apple-shaped capability.

**Meta Llama** —  $\Lambda = 0.0104$ , **Sub-Critical**. Landscape leader under v1.0 min()  $\beta$ . The v2.0 geometric mean reveals fragility. Llama 4 EU geofencing and the Avocado closed-source transition signal worsening structural position.

**Mistral/DeepSeek** —  $\Lambda = 0.0314$ , **Approaching Critical**. The only pattern above 0.03. Balanced incentive structure. Structural vulnerability is geopolitical. The next 18 months will determine whether geopolitical pressure or tooling maturation moves faster.

**The Autonomaton** —  $\Lambda = 0.0001$ , **Structurally Inert**. Highest structural base in the landscape ( $S \times R = 0.77$ ).  $V = 0.2$  reflects pre-publication status.  $\beta = 6.3$  reflects the absence of an exogenous forcing function. The score does not flatter the sponsoring organization.

## The operator's brief

**1. Stop renting your cognition.** Centralized APIs are not buying capability. They are renting logic the vendor retires on the vendor's schedule. Build against abstractions that outlive any single model.

**2. Freeze the per-seat AI upgrade.** 3.3% paid conversion on Copilot across 450M seats is what happens when AI is layered onto rigid legacy interfaces. If the AI seat does not produce defensible return, it does not belong in the renewal.

**3. Pivot toward sovereignty before the economics force it.** At 2M+ daily tokens, self-hosting open-weight models produces 80–90% savings against centralized APIs. *The organizations that deploy specialized engineering talent now will be 18 months ahead of peers who wait for the economics to force the move.*

**4. Build the middleware, not the model commitment.** The CTO's architectural responsibility is no longer to pick the winning model — it is to build middleware that makes any given model replaceable.

## Conflict of interest

The Grove Foundation publishes this framework and champions the Autonomaton architecture. The Autonomaton scores  $\Lambda = 0.0001$  — the lowest in the landscape. If Grove were tilting the methodology, that score would not appear in the table. Every input variable and sub-score is cited to source. The methodology is CC BY 4.0. Independent practitioners are encouraged to re-score the landscape and publish the differences.

## About this research

This is the first quarterly  $\Lambda$  landscape audit. Ninety-six sources. Eight patterns. Four historical calibrations. Full methodology disclosure at [the-grove.ai/standards](https://the-grove.ai/standards).

The next issue examines a contractual blind spot. AI vendor contracts signed in the summer of 2026 without a fulsome definition of telemetry are not subscriptions. They are permanent transfers of behavioral intelligence to the vendor that captured it under terminology that was never nailed down.

---

### The Grove Foundation

*"Design is philosophy expressed through constraint."*

CC BY 4.0 · March 2026 · Issue 002 · Date of Analysis: March 8, 2026

96 sources · 8 patterns · 4 historical calibrations